Tractable Pathfinding for the Stochastic On-Time Arrival Problem

Mehrdad Niknami
mniknami@berkeley.edu *

Samitha Samaranayake
samitha@berkeley.edu †

Alexandre Bayen
bayen@berkeley.edu ‡

arXiv:1408.4490v1 [cs.DS] 19 Aug 2014

Abstract

We present a new technique for fast computation of the route that maximizes the probability of on-time arrival in stochastic networks, also known as the path-based stochastic on-time arrival (SOTA) problem. We utilize the solution to the policy-based SOTA problem, which is of pseudopolynomial time complexity in the time budget of the journey, as a heuristic for efficiently computing the optimal path. We also introduce Arc-Potentials, an extension to the Arc-Flags pre-processing algorithm, which improves the efficiency of the graph pre-processing and reduces the computation time. Finally, we present extensive numerical results demonstrating the effectiveness of our algorithm and observe that its running time when given the policy (which can be efficiently obtained using pre-processing) is almost always linear in the length of the optimal path for our test networks.

*Electrical Engineering and Computer Science, University of California, Berkeley.
†Systems Engineering, University of California, Berkeley.
‡Electrical Engineering and Computer Science, University of California, Berkeley.

1 Introduction

In stochastic networks, a fast path in expectation is not necessarily a reliable one. High variability in traffic conditions can cause significant delays in the duration of a trip; therefore, it may be beneficial to allow for slightly longer paths in exchange for much stronger guarantees on reliability.

The *stochastic on-time arrival* problem [1], or SOTA for short, concerns itself with this trade-off. In the SOTA problem, one seeks to maximize the probability of on-time arrival at a destination d upon departing from a source s with a given travel time budget T in the face of uncertainty in the travel time. The travel time across each link is assumed to be a random variable with some arbitrary probability distribution $p(\cdot)$.

Unsurprisingly, the SOTA problem has received considerable attention in transportation engineering in the past decade, and has the potential for numerous applications in reliable routing, such as for on-time arrival at an airport, prompt delivery of perishable goods, and the like. Furthermore, the SOTA formulation can be extended to more general objectives by simply choosing a different utility function to the probability of on-time arrival [2]. The solution methods for the SOTA problem can still be applied to these generalized utility functions without requiring any modification, further motivating the search for efficient solutions.

Formulation. There are two primary variants of the SOTA problem, each with their own set of applications. The path-based SOTA problem, which is also referred to as the *shortest path problem with on-time arrival reliability* [3], consists of finding the most reliable fixed path to the destination. The policy-based SOTA problem, on the other hand, consists of computing a *routing policy*—rather than a fixed path—such that, at every intersection, the choice of the next direction depends on the current state (remaining time budget).[1]

The policy-based approach is preferred when the user has a device such as a GPS receiver available for receiving directions, as it is possible for the system to adapt the path to the remaining time budget, potentially allowing for achieving a higher likelihood of on-time arrival in the process. In some settings, however, such a device with sufficient power and positioning signal may not be available, or the passenger may desire to know the entire trip prior to departure. In such a scenario, one desires to choose most reliable path a priori.

The policy-based SOTA problem is generally solved in discrete-time, as it allows for convolutions of arbitrary distributions. Furthermore, computation of the policy involves a subsequent maximization step, which mixes the distributions and prevents one from finding closed-form solutions in continuous-time, even for distributions that are closed under convolution. Fan and Nie [5] presented a successive-approximation algorithm for the policy-based discrete SOTA problem, which was subsequently improved by Samaranayake et al. [4] to a dynamic-programming algorithm that solves the problem in pseudopolynomial-time utilizing a sequence of speedup techniques including Zero-Delay Convolution [6, 7]. Sabran et al. [8] further reduced query times by extending the graph pre-processing algorithms Reach [9] and Arc-Flags [10] to the SOTA problem.

There is no known polynomial time solution for the general path-based SOTA problem [3]. Nikolova et al. [11, 12], Lim et al. [13] presented a polylogarithmic-time solution for special scenarios, such as when travel times have a Gaussian distribution. Nie and Wu [3] attempt to find an efficient heuristic algorithm for solving the problem using first-order stochastic dominance of paths, allowing for more general probabilistic models. However, this algorithm is not fast enough for practical application in real-world networks.

[1] In this paper, we only consider time-invariant travel-time distributions. The problem formulation can be extended to incorporate time-varying distributions as discussed in [4].

Contribution. In this paper, we introduce a new algorithm for the path-based SOTA problem. We utilize the solution to the policy-based SOTA problem as a heuristic for guiding the path search toward the destination.[2] As this requires computation of the underlying SOTA policy, we also present an improvement to an existing pruning technique known as Arc-Flags to help speed up the policy computation. We then present a simple and efficient pre-processing approach that, to the best of our knowledge, is the first attempt at pre-processing for the path-based SOTA problem. We finally present our results, and show that our heuristic almost always finds the optimal path with computation time *linear* in the length of the path, when the policy has already been computed. The efficiency of the solution is due to the high accuracy of the SOTA policy in directing the search for the optimal path.

2 Preliminaries

We are given a stochastic network in the form of a directed graph $G = (V, E)$ where each edge $(i, j) \in E$ has an associated probability distribution $p_{ij}(\cdot)$ representing the travel time across that edge.[3] Let $s \in V$ be the source node, $d \in V$ be the destination, and $T \in \mathbb{R}^+$ be the remaining time budget at the source.

Definition 1 (Minimum Travel Time) Define δ_{ij} to be the *minimum travel time* along the edge from node i to node j:

$$\delta_{ij} = \min_t \{t \ : \ p_{ij}(t) > 0\}$$

Definition 2 (On-Time Arrival Probability) Let $u_{ij}(t)$ be the arrival probability at the destination d when leaving leaving node i with a travel-time budget of t, first traversing edge $(i, j) \in E$, and subsequently following the optimal policy. Then the on-time arrival probability $u_i(t)$ and the policy (optimal subsequent node) $w_i(t)$ at node i, can be defined in a dynamic-programming fashion as follows:

$$\begin{aligned}
u_{ij}(t) &= \int_{\delta_{ij}}^{t} u_j(t-\tau) p_{ij}(\tau)\, d\tau \\
u_i(t) &= \max_{j:\,(i,j)\in E} u_{ij}(t) && \text{(on-time arrival probability)} \\
w_i(t) &= \arg\max_{j:\,(i,j)\in E} u_{ij}(t) && \text{(optimal policy for intermediate nodes)} \\
u_d(\cdot) &= 1 && \text{(initial condition for destination)}
\end{aligned}$$

This dynamic-programming formulation of the problem results in the set of ordering constraints in Lemma 1 that must be satisfied for correct computation of the solution.

Lemma 1 (Ordering Constraints and Optimal Ordering) *Computing $u_i(t)$ requires first computing $u_j(\tau)$ $\forall\, \tau$ where*

$$0 \le \tau \le t - \delta_{ij} \quad \forall (i, j) \in E$$

Furthermore, through repeatedly applying these constraints in a Dijkstra-like fashion, one can obtain the minimal sequence of updates that need to performed to solve the policy-based SOTA problem.

[2] This essentially results in an application of the A^* algorithm to the path-based SOTA problem, using the policy-based SOTA solution as an admissible heuristic.

[3] For simplicity and without loss of generality we assume at most one edge exists between any pair of nodes in each direction.

PROOF The correctness of the constraint follows from the definition of δ_{ij}: $p_{ij}(\tau) = 0$ for all $\tau < \delta_{ij}$. Hence $u_j(t-\tau)$ cannot affect the product $u_j(t-\tau)p_{ij}(\tau)$ and is therefore not required for computing $u_i(t)$. The proof that the sequence of updates is minimal can be found in [4].

For notational simplicity, we present the SOTA problem in continuous-time throughout this paper, with the understanding that the algorithms will be applied in discrete time in intervals of Δt where $\Delta t \leq \delta_{ij}\ \forall (i,j) \in E$. The discretization is straightforward, but a discussion of the procedure can be found in [4].

Remark 1 (Policy Computation) Computation of the policy $u_s(T)$ for a source s and time budget T using direct convolution takes $O(|E|t^2)$ time, as computing $u_s(T)$ could in the worst case require convolutions of length $O(T)$ for each edge in the graph. (A more detailed explanation can be found in [4].) However, using a stochastic dynamic programming technique known as *zero-delay convolution* (ZDC) [6] in lieu of direct convolution, the time complexity can be reduced to $O(|E|t\log^2 t)$. Details on how to apply ZDC to the SOTA problem can be found in [6, 7].

3 Path-Based SOTA

In the deterministic setting, the reason for the efficiency of Dijkstra's algorithm can be attributed to *sub-path optimality*: namely, any optimal route to a destination node d that includes some intermediate node i necessarily includes the optimal path from i to d. Unfortunately, as illustrated in Appendix A, this does not hold in the stochastic setting. We therefore seek an efficient algorithm for the computing the SOTA path that searches the state space intelligently.

3.1 The SOTA Policy as a SOTA Path Heuristic

Path $\preceq$ Policy. It is obvious that a predetermined path cannot guarantee a higher arrival probability than a policy, as the policy would otherwise simply direct the user to use the same path irrespective of the realized travel time. Therefore, the on-time arrival probability of the most reliable policy provides a natural upper bound on the reliability of the most reliable path, which we will use as the foundation of our new algorithm.

Upper Bound. Consider a fixed path P from the source s to node i. Let $q_{si}^P(t)$ be the travel time distribution along P from node s to node i, i.e., the convolution of the travel time distributions of every edge in P. Upon arriving at node i at time t, let the user follow the optimal policy toward d, therefore reaching d from s with probability density $q_{si}^P(t)u_i(T-t)$.

The reliability of following path P to node i and subsequently following the optimal policy toward the destination d is:[4]

$$r_{si}^P(T) = \int_0^T q_{si}^P(t)u_i(T-t)\,dt$$

We will use this mixed formulation as a continually improving upper bound on the reliability to direct the SOTA path search in an efficient manner.

3.2 Algorithm for Computing the SOTA Path from the SOTA Policy

The simple approach for computing the SOTA path would be to enumerate through all possible paths by exploring the path tree until we reach the destination d. However, this is clearly intractable and we need to explore the path tree in an intelligent manner. We will use the SOTA policy as a heuristic to guide the search through the exponential path space efficiently and exploit the fact

[4] The bounds of this integral can be slightly tightened through inclusion of the minimum travel times, but this has been omitted for simplicity.

that the policy solution is an admissible upper bound on the path's reliability to terminate the search. Starting at the origin with $r^P_{ss}(T)$, we add the mixed (path and policy) solutions $r^{P_i}_{si}(T)$ for all the neighbors i of s to a priority queue. We then pick the mixed solution with the highest upper bound and continue the process. Given a mixed solution $r^P_{si}(T)$, since the policy heuristic is admissible, we know that the reliability of the solution can never improve when extending the path by picking an edge at the path-policy interface. Thus, if we start at the origin and keep extending paths based on this heuristic, once we have a path terminating at the destination, *we are certain that this path is the most reliable.* We detail the steps of this procedure in Algorithm 1.

Algorithm 1 Algorithm for computing the optimal SOTA path

Notation: $*$ is the convolution operator and $\|$ is the concatenation operator
for all $i \in V$, $0 \leq t \leq T$ **do**
    Compute[5] $u_i(t)$
$Q \leftarrow$ PriorityQueue()
Q.Push($u_s(T)$, ([1.0], [s])) ▷ Push initial path at source keyed with its reliability
while Q is not empty **do**
    $(r, (q, P)) \leftarrow Q$.PopMax() ▷ Extract most reliable path so far
    $i \leftarrow P[|P| - 1]$ ▷ Get the last node in the path
    if $i = d$ **then** ▷ Check if path terminates at destination
        return P ▷ Return solution (or fall through to find the next-best path)
    for all $j \in E$.Neighbors(i) **do**
        if $j \notin P$ **then** ▷ Skip paths with loops
            Q.Push($(q * u_j)[T]$, $(q * E.\text{Cost}(i, j), P \parallel [j])$) ▷ Extend the old path
return nil ▷ No path found

3.3 Analysis

The practical efficiency of our algorithm depends crucially on the travel time distributions along each road. The worst-case scenario for the algorithm occurs when the policy directs the traveler toward the worst possible path, which may happen in pathological cases where the ability to delay decisions in the journey significantly increases the flexibility and benefit of choosing a route that reaches the destination on time. Such a scenario, or even one in which the policy frequently chooses a suboptimal path would result in an exponential time complexity for our algorithm.

However, it is difficult to imagine such a scenario happening in practice, and we conjecture that when travel time distributions are time-invariant, such an event would be extremely unlikely to occur in real-world road networks. As will be shown later, experimentally, we came across very few cases in which the path computation time dominated the policy computation time, and even in those cases, they were still quite reasonable and extremely far from such a worst-case scenario. An interesting open problem is to define a set of characteristics for the network and the probability distributions that guarantee polynomial run-times, similar in nature to the *Highway Dimension* result in [14].

4 Pre-processing with Arc-Potentials

In standard deterministic pathfinding, pre-processing techniques such as *Arc-Flags* [10], *reach-based routing* [9, 15], *contraction hierarchies* [16], and *transit node routing* [17] have been developed that can help decrease the query time by orders of magnitude. Some of these approaches allow for

[5] Can be limited to those i and t reachable from s in time T, and can be further sped up through existing policy pre-processing techniques such as Arc-Flags.

pruning the search space based solely on the destination node, while others also take the source node into account, allowing for better pruning at the cost of additional pre-processing.

The structure of the SOTA problem makes it difficult to apply such techniques to the computation of the policy. Previously, Arc-Flags and Reach have been successfully adapted to the policy-based SOTA problem in [8], resulting in stochastic Arc-Flags and Directed Reach. While at first glance one may be tempted to directly apply these algorithms to the computation of the policy, a naive application of source-dependent pruning techniques (such as Directed Reach or source-based Arc-Flags) would result in an incorrect solution, as the intermediate policy computations form separate SOTA problems which do *not* all share the same starting regions. Thus our options for policy pre-processing are limited to destination-based (source-independent) techniques such as stochastic Arc-Flags.

While at first disappointing, this limitation of source-based approaches can be overcome through *path-based* pre-processing; that is, pruning the graph to include only those edges which may be part of the most reliable path. This will generally result in a substantial reduction of the search space on which the policy must be computed.[6] However, naively re-computing a new path for the entire range of time budgets up to our desired target is intractable and therefore calls for a more efficient solution.

In this section, we first present an improvement upon Arc-Flags (which we call *Arc-Potentials*) to increase the efficiency of the pre-processing algorithm in the general case, and we subsequently describe an efficient procedure for computing the path for a large range of time budgets.

4.1 Arc-Potentials

Consider partitioning the graph G into r regions. In the deterministic setting, Arc-Flags allow us to pre-process and prune the search space as follows. For every arc (or edge) $(i, j) \in E$, we define a bit-vector of length r that denotes whether or not this arc is on an optimal policy ending at some node in region r. We then pre-compute these Arc-Flags, and store them for later pruning of the graph at query time. Furthermore, deterministic Arc-Flags can be extended to the dynamic setting [18], in which the flags are updated with low recomputation cost after the road network is changed.

Sabran et al. [8] applied Arc-Flags to the policy-based SOTA problem as follows: each bit vector was defined to represent whether or not its associated arc is *on an optimal policy* to some destination in the target region associated with each bit. The problem with this approach, however, is that it requires computing arc-flags for *multiple* target budgets, each of which themselves take a considerable amount of space. Instead, we propose more efficient alternative, which we call *Arc-Potentials*:

Definition 3 (Arc-Potentials) For a given destination region D, we define the *arc activation potential* ϕ_{ij} of the edge from node i to node j to be the *minimum* time budget at which the arc becomes part of an optimal policy to some destination $d \in D$.[7]

Once the set of realizable edges under the optimal policy is known, computing the activation potentials ϕ_{ij} for a static graph is a straightforward process, similar to that of Arc-Flags. As with Arc-Flags, to keep the space complexity linearithmic in the size of the network, one can choose $|D|$ to be tunable constants approximately proportional to $\log |E|$, depending on the level of accuracy and storage desired.

[6] It is of no concern that the policy on the pruned graph differs from that on the original network, as any edges which are known not to belong to the optimal path can by definition be entirely eliminated from consideration for computing the optimal path.

[7] See Appendix B for a generalized definition of Arc-Potentials with better pruning.

It is not, however, immediately obvious how one should compute the set of realizable edges, as only a subset of the physically reachable edges for a given time budget can be reached when following the optimal policy. We therefore provide an algorithm for computing the set of realizable edges given in Algorithm 2, based on the definition of realizability.

Definition 4 (Realizability) Let $f_j(t) = 1$ if node j is reachable at time t while following an optimal policy, and 0 otherwise. Formally, if $\mathbf{1}$ is the indicator function, we have:

$$f_j(t) = \mathbf{1}(j = s \ \vee \ \exists\, i \in V, \tau \in \mathbb{R} : \ w_i(t+\tau) = j \ \wedge \ u_i(t+\tau) p_{ij}(\tau) > 0)$$

Algorithm 2 shows how the realizable edges and corresponding Arc-Potentials are computed. Note that $\chi[k] = (i, (t^-, t^+))$ refers to the update-order array required for computation of the SOTA policy, in which $i \in V$ is some node for which $u_i(t)$ is to be computed for $t^- \leq t \leq t^+$. Note that χ can be any topological ordering of the updates that satisfies the equation in Lemma 1, including optimal-ordering used in [4] for the computation of the SOTA policy.

Algorithm 2 Algorithm for Computing Arc-Potentials

First, we compute the set of realizable edges
χ is an update order array of the form $\chi[k] = (i, (t^-, t^+))$
ϕ_{ij} is the minimum time budget for which edge (i, j) is realizable for some policy toward a $d \in D$
for all $i \in V$ **do** ▷ Compute optimal policy from all nodes
  Compute $w_i(t) \ \ \forall\ 0 \leq t \leq T$
for all $i \in V$ **do**
  $f_i(t) \leftarrow$ **false** $\ \forall\ 0 \leq t \leq T$
$f_s(t) \leftarrow$ **true** $\ \forall\ 0 \leq t \leq T$ ▷ Mark the source as reachable
Compute an update order χ from s to d
for $k = |\chi|$ to 1 **do** ▷ Iterate in reverse order (assuming $v_0 = s$, $t_0^+ = T$)
  $(i, (t^-, t^+)) \leftarrow \chi[k]$
  for all $t^- \leq t \leq t^+$ **do**
    $j \leftarrow w_i(t)$
    for all $\delta_{ij} \leq \tau \leq t$ **do**
      $f_j(t-\tau) \leftarrow f_j(t-\tau) \vee (f_i(t) \wedge (p_{ij}(\tau) > 0))$
Define $\phi_{ij} \leftarrow \infty \ \ \forall (i, j) \in E$
for all $d \in D$ **do**
  for all $(i, j) \in E$ **do**
    $\phi_{ij} \leftarrow \min(\phi_{ij}, \min\{t : f_i(t) = \textbf{true}\})$

Proposition 1 (Realizable Edges) *Algorithm 2 marks exactly the set of edges realizable under the optimal policy in $O(|V|T^2)$ time.*

Proof At every iteration, the algorithm selects the next node to update i, and for its optimal neighbor j at each time step marks all the times at which one could possibly arrive at j based on the probability distribution $p_{ij}(\cdot)$ and f_i. Since the node updates follow a topological ordering with respect to the update interval, no other subsequent update could later influence $f_i(t^- ... t^+)$, and thus the algorithm computes f_j correctly for all $t \in [t^-, t^+]$ for which the edge (i, j) is realizable. Since the time budget for every update in the update array can be on the order of T, and since each such time step requires marking $O(T)$ bits for each neighbor j of node i in this update, then the algorithm takes $O(|V|T^2)$ time, where the number of updated nodes is assumed to be $O(|V|)$.

Proposition 2 (Arc-Potentials) *Algorithm 2 correctly computes the activation potential of every arc and can be later used at query time to prune arcs that will not belong to the optimal policy for a given destination node and budget.*

PROOF By construction, at the end of the algorithm, ϕ_{ij} is less than or equal to the minimum time budget at which the edge from i to j becomes optimal toward any node in the destination region. In other words, it is guaranteed that edge $(i, j) \in E$ is not part of any optimal policy toward any destination $d \in D$ for budgets less than ϕ_{ij}. Therefore, the edge can be pruned for these budgets.

While this algorithm has quadratic time complexity in the time-budget ($O(|V|T^2)$), it is still fast enough and sufficient for the time budgets and discretization intervals of interest for our networks, in part because of the practical efficiency of bit-vector manipulation. However, on larger networks or for smaller Δt, the time complexity can be reduced to $O(|V|T \log^2 T)$ (which is then dominated by the policy computation time) via transforming f into a real number and then performing zero-delay convolution from the source to the destination, substituting addition and multiplication for the Boolean operators as appropriate.

4.2 Efficient Path-based Pre-processing

Path-based pre-processing requires finding all edges that belong to optimal paths up to a desired time budget T. While a naive approach requires $O(T)$ times the work of finding a single path, this would be quite inefficient in practice, as the optimal path rarely changes as the time budget is increased. Not surprisingly, it is possible to improve this method significantly.

To improve the efficiency of the naive approach, we first observe that algorithm 1 requires computing the policy reliability vector $u_i(\cdot)$ the travel-time distribution $q^P_{si}(\cdot)$ for each path under consideration, and the inner product of the two vectors, which are *shifted* relative to each other as the time budget is progressively increased. Computing $u_i(\cdot)$ for all time budgets of interest is already a simple task, as the dynamic-programming algorithm allows us to pre-compute it for all nodes before any paths are investigated. We also note that the travel-time distribution for a given path P is *not* a function of the remaining time budget[8], and thus extending $q^P_{si}(\cdot)$ to include an additional edge $(i, j) \in E$ amounts to a simple $O(T)$-time convolution of the two distributions, since both are already stored in memory. Therefore, the only remaining task is to compute $O(|P|)$ inner products (one for every node along the path) of size $O(T)$, one for each time budget T under consideration, which can be further accelerated via reformulating the computation of the inner products as a simple zero-delay convolution of the two distributions.

5 Experimental Results

We tested the performance of our C++ implementation of the pathfinding algorithm along with Arc-Potentials on two networks. The first is a small San Francisco network with 2643 nodes and 6588 edges for which real-world travel-time distribution data was available as a Gaussian mixture model [19], and the second is a relatively large Luxembourg network with 30647 nodes and 71655 edges for which travel-time distributions were synthesized from road speed limits as real-world data was not available. San Francisco pre-processing and queries were run on a quad-core Intel Core i7 4700MQ machine with 6 GiB of memory, and Luxembourg pre-processing and queries were run on 18× 1.9 GHz AMD Opteron(tm) 6168 processors with 30 GiB of shared memory.

To generate random problem instances, we picked two nodes uniformly at random from the graph, and found the LET path (that is, the path with the least expected travel time) between

[8] For efficiency, it may be desireable to truncate this vector to the maximum time budget under consideration.

them. We then chose a budget uniformly at random from the 5^{th} to 95^{th} percentile of possible travel times[9]. All the results use Zero-Delay Convolution for computing the policy, with $\Delta t = 1$ second.

Figures 1a and 1b compare the effectiveness of applying Arc-Potentials to the underlying policy (considering only the destination region) with that of applying Arc-Potentials to the path (considering both the source and the destination region). As expected, path-based pre-processing performs much better than purely policy-based pre-processing, and both approaches become better as the region sizes are decreased. Nevertheless, the incremental benefit of using smaller regions is low; the majority of the search space reduction is achieved via a small number of regions, which shows that the algorithm is extremely effective despite its low space requirement.

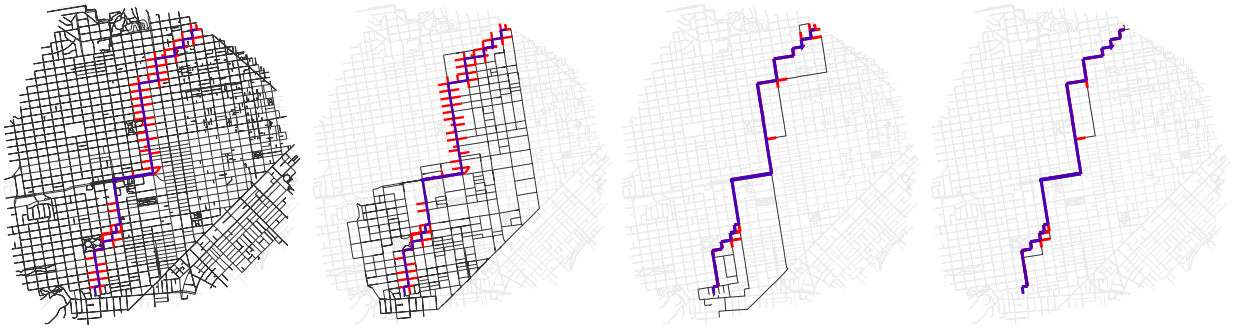

0 regions (unpruned) 13^2 regions (policy-based) 10^2 regions (path-based) Maximum (path-based)

(a) A San Francisco instance with a budget of 837 seconds from the source (top) to the destination (bottom).

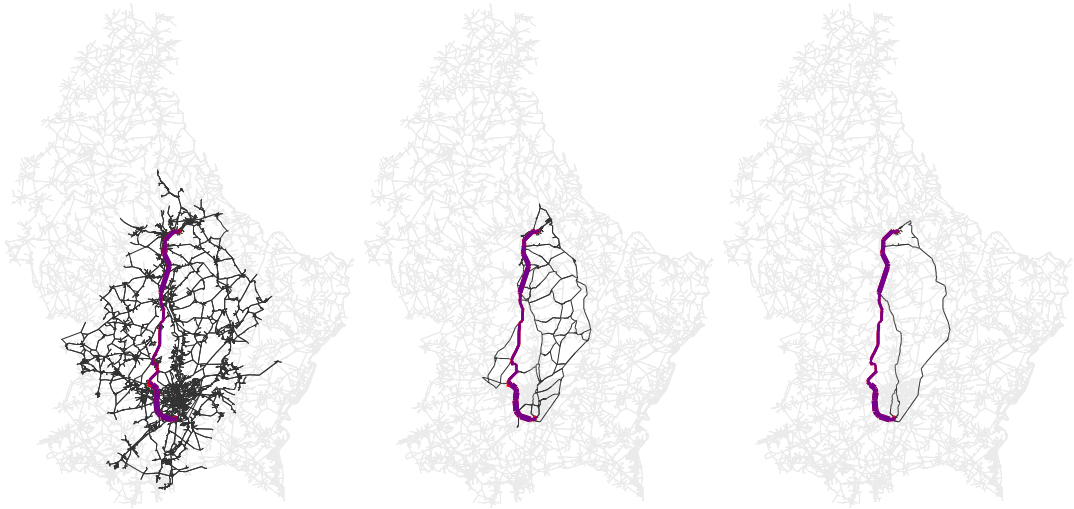

0 regions (unpruned) 23^2 regions (policy pruning) 23^2 regions (path pruning)

(b) A Luxembourg instance with a budget of 3165 seconds from the source (bottom) to the destination (top).

Figure 1: The SOTA solution for random instances of the San Francisco and Luxembourg networks. Light gray edges were pruned from the graph. Blue edges constitute the optimal path to the destination. Red edges represent suboptimal edges belonging to paths on the queue at the termination of the algorithm.

[9] Time budgets which are too high (i.e. $\approx$ 100% reliability) are uninteresting because the LET path already provides full reliability. Time budgets which are too low ($\approx$ 0% reliability) are also not of interest, since they offer almost zero probability of success.

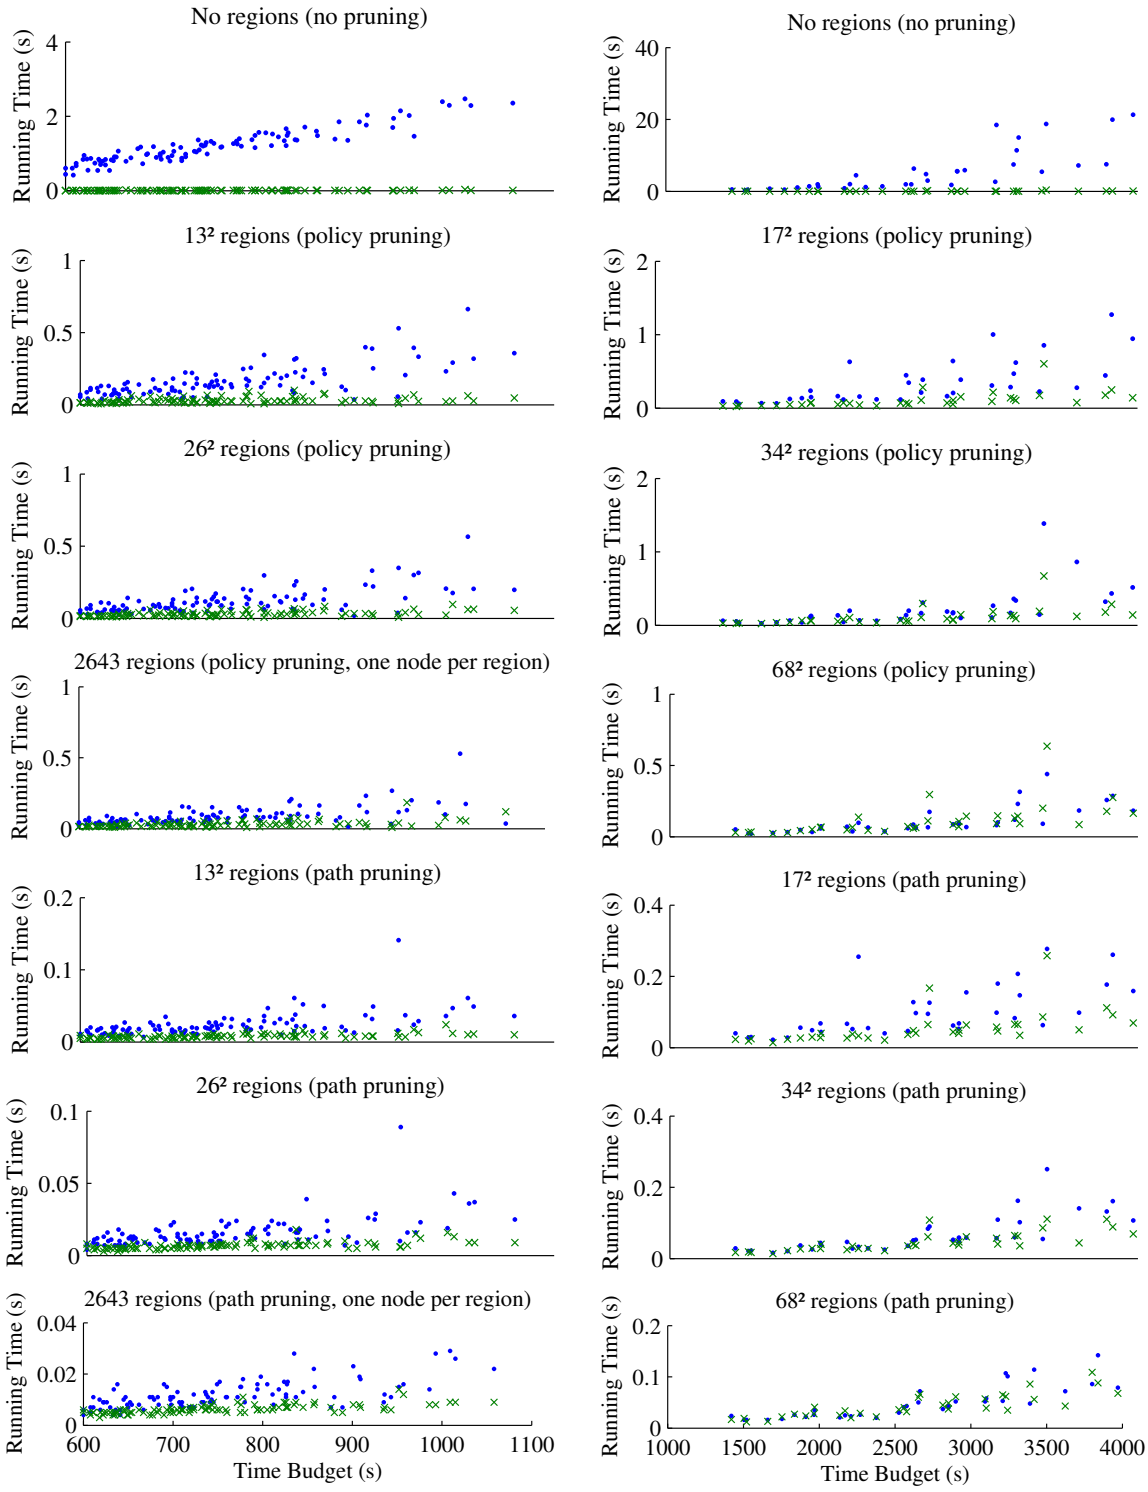
No regions (no pruning)
13² regions (policy pruning)
26² regions (policy pruning)
2643 regions (policy pruning, one node per region)
13² regions (path pruning)
26² regions (path pruning)
2643 regions (path pruning, one node per region)
17² regions (policy pruning)
34² regions (policy pruning)
68² regions (policy pruning)
17² regions (path pruning)
34² regions (path pruning)
68² regions (path pruning)
Running Time (s)
Time Budget (s)

(a) San Francisco (100 random instances)

(b) Luxembourg (30 random instances)

Figure 2: The running time of our algorithm plotted against the time budget for each problem instance. Sources and destinations were selected randomly and the network was pre-processed for these regions to compare the running times against the unpruned graph.[10] Blue dots represent computation of the policy, and green ×'s represent computation of the path using that policy.

Figure 2 shows the running times of the policy and path-based SOTA algorithms. We can see that, in practice, computation of the optimal path (when the policy is known) takes less time than computing the policy itself, resulting in extremely fast queries in practice.

Finally, figure 3 shows how the algorithm's running time scales with the length of the optimal path for the San Francisco network, both in terms of the number of edges in the optimal path as well as in the mean travel time of the path.

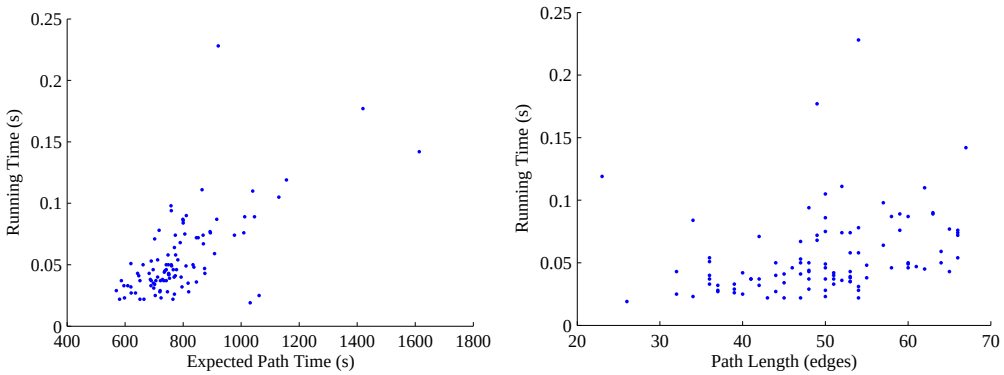

Figure 3: Running time of the pathfinding algorithm on random instantiations of the San Francisco network compared to the mean travel time along the optimal path (left) and the number of edges in the path (right). We can see the running time is linear in most cases.

These numerical results show that our pathfinding algorithm is likely to work extremely well in real-world scenarios, especially when combined with pre-processing.

6 Conclusion and Future Work

We have presented a novel algorithm for computing the path that maximizes the on-time arrival probability to a destination based on first computing the optimal policy for arriving on-time at the destination. We have also presented algorithms to speed up computation of the underlying policy, including a generalization of Arc-Flags pre-processing algorithm that we call Arc-Potentials. We have furthermore applied and implemented these algorithms to sample transportation networks for moderate-sized cities, and demonstrated their potential for high efficiency in real-world networks.

While unobserved in practice, there remains the possibility that our algorithm may perform poorly on stochastic networks in which the optimal policy is a poor heuristic for the path reliability. Proofs in this direction have remained elusive, and investigation of whether such scenarios can occur in realistic networks remains an important theoretical question of great interest for future research.

In the absence of theoretical improvements, however, our empirical algorithm provides an extremely efficient practical alternative to the state-of-the-art techniques for solving the path-based SOTA problem. It would be of interest to investigate whether our policy-based approach can be combined with other pathfinding techniques to gain additional speedup, including the stochastic dominance method of Nie and Wu [3] as well as generalizations of more traditional pathfinding techniques, such as bidirectional search. It is likely that combining several such approaches can provide even more speedup for commercial applications of the SOTA problem.

[10]Resource limitations constrained the number of benchmarks we were able to run on the synthetic Luxembourg network.

Appendix A Stochastic Sub-Path Optimality

In deterministic pathfinding, we may safely construct an optimal path by combing optimal sub-paths. For example, if the optimal path from s to d includes node i, the path can be constructed by combining the optimal paths from s to i and i to d. However, this property does not hold in the stochastic setting; even when a edge or sub-path is strictly dominated by others, we may not be able to prune it from the graph without compromising the correctness of the solution.

A simple counterexample illustrates this fact. Consider the example in figure 5, in which the objective is to find the most reliable path from v_1 to v_3 with a time budget of 4. Since the edge e_4 takes at least 1 time unit, any feasible path from v_1 to v_2 must take 1, 2, or 3 time units. As edge e_2 is dominated by edges e_1 and e_3 for these time budgets, one might be tempted to ignore this edge when searching for the optimal path. However, this would be wrong, as the optimal path is (e_2, e_4). Thus, unlike in the deterministic setting, we may not necessarily prune edges based on sub-path optimality.

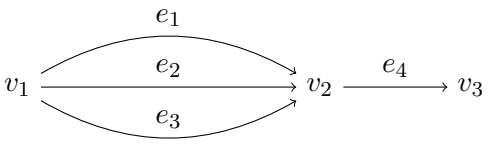

Edge	$\mathbb{P}(t \leq 1)$	$\mathbb{P}(t \leq 2)$	$\mathbb{P}(t \leq 3)$	$\mathbb{P}(t \leq 4)$
e_1	0	0.9	1	1
e_2	0.5	0.8	0.9	1
e_3	0.6	0.6	0.6	1
e_4	0	0.5	1	1

Figure 5: A counterexample showing that sub-path optimality does not hold in the SOTA setting.

Appendix B Asymptotically Exact Pruning with Arc-Potentials

The pruning provided by Algorithm 2 is suboptimal because it only stores the *"activation"* potential of every edge. This implies that, for large budgets, every edge is potentially active and therefore no pruning can occur.

Generalized Arc-Potentials. A simple generalization of the Arc-Potentials pruning algorithm can easily allow for asymptotically exact pruning at relatively low cost: we store the actual potential **intervals** during which the arc is active, rather than merely the lower bound. In other words, one would keep a list of the k lowermost intervals during which the arc is active, along with an additional lower bound on the next activation potential.

The storage space for this generalization clearly increases linearly in k, and the running time essentially increases only to the extent that T is desired to be increased. Furthermore, depending on the particular probability distributions involved, the maximum arc (re-)activation potential can in fact be bounded for some arcs. In such cases, it may be very well desirable to compute all activation intervals entirely and store them for later retrieval, if the storage and computation costs are not prohibitive. However, in our experiments this was deemed unnecessary as Arc-Potentials were already sufficient for significant pruning in the time budgets of interest in our networks.

References

[1] Yueyue Fan, Robert Kalaba, and JE Moore II. Arriving on time. *Journal of Optimization Theory and Applications*, 127(3):497–513, 2005.

[2] Arthur Flajolet, Sebastien Blandin, and Patrick Jaillet. Optimal routing on stochastic networks with transportation applications. SIAM Conference on Optimization, 2014.

[3] Yu Marco Nie and Xing Wu. Shortest path problem considering on-time arrival probability. *Transportation Research Part B: Methodological*, 43(6):597–613, 2009.

[4] Samitha Samaranayake, Sebastien Blandin, and Alexandre Bayen. A tractable class of algorithms for reliable routing in stochastic networks. *Transportation Research Part C: Emerging Technologies*, 20(1): 199–217, 2012.

[5] Yueyue Fan and Yu Nie. Optimal routing for maximizing travel time reliability. *Networks and Spatial Economics*, 6(3-4):333–344, 2006.

[6] Brian C Dean. Speeding up stochastic dynamic programming with zero-delay convolution. *Algorithmic Operations Research*, 5(2):Pages–96, 2010.

[7] Samitha Samaranayake, Sebastien Blandin, and Alexandre Bayen. Speedup techniques for the stochastic on-time arrival problem. In *OASIcs-OpenAccess Series in Informatics*, volume 25. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2012.

[8] Guillaume Sabran, Samitha Samaranayake, and Alexandre Bayen. Precomputation techniques for the stochastic on-time arrival problem. In *ALENEX*, pages 138–146. SIAM, 2014.

[9] Ronald Gutman. Reach-based routing: A new approach to shortest path algorithms optimized for road networks. In *ALENEX/ANALC*, pages 100–111, 2004.

[10] Moritz Hilger, Ekkehard Köhler, Rolf Möhring, and Heiko Schilling. Fast point-to-point shortest path computations with arc-flags. *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, 74:41–72, 2009.

[11] Evdokia Nikolova, Jonathan Kelner, Matthew Brand, and Michael Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *Algorithms–ESA 2006*, pages 552–563. Springer, 2006.

[12] Evdokia Nikolova, Matthew Brand, and David Karger. Optimal route planning under uncertainty. In *ICAPS*, volume 6, pages 131–141, 2006.

[13] Sejoon Lim, Christian Sommer, Evdokia Nikolova, and Daniela Rus. Practical route planning under delay uncertainty: Stochastic shortest path queries. *Robotics*, page 249, 2013.

[14] Ittai Abraham, Amos Fiat, Andrew Goldberg, and Renato Werneck. Highway dimension, shortest paths, and provably efficient algorithms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 782–793. Society for Industrial and Applied Mathematics, 2010.

[15] Andrew Goldberg, Haim Kaplan, and Renato Werneck. Reach for a^*: Efficient point-to-point shortest path algorithms. In *ALENEX*, volume 6, pages 129–143. SIAM, 2006.

[16] Robert Geisberger, Peter Sanders, Dominik Schultes, and Daniel Delling. Contraction hierarchies: Faster and simpler hierarchical routing in road networks. In *Experimental Algorithms*, pages 319–333. Springer, 2008.

[17] Holger Bast, Stefan Funke, and Domagoj Matijevic. Transit: ultrafast shortest-path queries with linear-time preprocessing. *9th DIMACS Implementation Challenge [1]*, 2006.

[18] Gianlorenzo DAngelo, Daniele Frigioni, and Camillo Vitale. Dynamic arc-flags in road networks. In *Experimental Algorithms*, pages 88–99. Springer, 2011.

[19] Timothy Hunter, Pieter Abbeel, and Alexandre Bayen. The path inference filter: model-based low-latency map matching of probe vehicle data. In *Algorithmic Foundations of Robotics X*, pages 591–607. Springer, 2013.